An analysis of budgeted parallel search on conditional Galton-Watson trees

David Avis[1,2] and Luc Devroye[1]

[1]School of Computer Science, McGill University, Montréal, Québec, Canada
[2]School of Informatics, Kyoto University, Kyoto, Japan

April 3, 2017

arXiv:1703.10731v1 [cs.DS] 31 Mar 2017

Abstract

Recently Avis and Jordan have demonstrated the efficiency of a simple technique called budgeting for the parallelization of a number of tree search algorithms. The idea is to limit the amount of work that a processor performs before it terminates its search and returns any unexplored nodes to a master process. This limit is set by a critical budget parameter which determines the overhead of the process. In this paper we study the behaviour of the budget parameter on conditional Galton-Watson trees obtaining asymptotically tight bounds on this overhead. We present empirical results to show that this bound is surprisingly accurate in practice.

Keywords: Parallel tree search, random Galton-Watson tree, probabilistic analysis of algorithms, branching process.

1 Introduction

The majority of algorithms have been designed, analyzed and implemented to run on single core processors. While multicore hardware is now ubiquitous these algorithms profit little, if any, from the additional processing power. On the other hand, the study of parallel algorithms also has a long history. Issues involved are complex and include architecture design, communication, data sharing, interrupts, deadlocks, load balancing, and the distinction between shared memory and distributed computing. For a comprehensive reference, see Mattson (2004). This complexity presents serious challenges to developing theoretical analyses to explain successful empirical results. Indeed many recent highly publicized computational success stories, such as Computer Go, Machine Learning and Deep Learning for AI, make massive use of parallel computation but have little or no theoretical foundation for the underlying algorithms.

The purpose of this paper is to make a modest contribution to this filling this gap by giving a theoretical analysis of a simple and effective scaleable parallelization technique, described by Avis and Jordan (2015, 2016), that is applicable to a certain class of tree search algorithms. Trees in this class must be definable by an oracle which, given any node in the tree, gives the children of that node (if any) in an arbitrary but fixed order. Examples of such trees are those used in reverse search, satisfiability testing, branch and bound, and games. In principle it is easy to parallelize searching trees in this class since subtrees can be explored independently by separate processors. The major problem is one of *load balancing* since in typical applications the trees to be explored are highly unbalanced. Processors assigned to large subtrees will be busy long after processors assigned to smaller subtrees are idle. The load balancing problem has been well studied, see for example Xu and Lau (1997).

Avis and Jordan approach the load balancing problem by using a simple technique called *budgeting*. A search of a subtree of the original tree terminates after a given number of nodes, called the *budget*, have been generated and returns all unexplored nodes to a master process which stores them on a job list. In a multiprocessor setting, the master process assigns nodes from the job list to available processors, called workers, that proceed in parallel, repopulating the list whenever their budget is reached. Using this technique a special purpose wrapper, *mplrs*, was written that made use of the largely unmodified *lrs* code. Experimental results showed near linear speedups when using up to several hundred processors. In the second paper the authors describe a generic common wrapper that can be used with a wide variety of legacy codes. They gave applications to other reverse search codes and, more recently, to satisfiability testers, again obtaining substantial speedups. The resulting software, called *mts*[1], is freely available. Budgeting may be seen as a very simple form of *job stealing*, as described by Blumofe and Leiserson (1999). However in their framework individual workers maintain their own job lists and poll each other as necessary to obtain additional work, requiring communications between workers, interrupt handling, deadlock avoidance and so on.

Budgeting has several very practical advantages. Firstly it does not require communication between workers and they do not need to process interrupts. This means that parallelism does not have to be added to the existing legacy code. Secondly large subtrees are automatically broken up, since they exceed the budget, whereas small subtrees are not, since they do not. Thirdly the master can dynamically control the size of the job list by varying the size of the budget when assigning work. Finally it is easy to checkpoint the process since all workers return to the master regularly due to the budget restriction. It is simply a matter of waiting for all workers to return and then outputting the remaining job list. On the other hand budgeting introduces two competing forms of overhead. One is the cost of restarting jobs on the master's list, which is directly proportional to the total size of the list. The other is the cost of idle workers if the list becomes empty. It is therefore very important to determine how the budget parameter affects the size of the job list.

The goal of this paper is to analyze the budgeting method when it is applied to Galton-Watson trees. In particular we study the critical budget parameter and how it determines the size and evolution of the job list. In Section 2 we formally introduce budgeting and give an explicit example of how it can be implemented in a depth first search setting. Next, in Section 3 we introduce the probabilistic model we will use for our analysis and state our main result. In Section 4 we describe a random walk which is closely related to the evolution of the list of unexplored nodes produced by budgeting. This is followed in Section 5 by a proof of our main result. Finally in Section 6 we make use of the *mts* framework to demonstrate the accuracy of our estimates on a sample of large Galton-Watson trees. We also give a discussion of how our main result can be used in a multiprocessor setting in order to choose an efficient budget parameter.

2 Budgeted tree search

We are given a tree T by its root, an upper bound Δ on the number of children of any node, and an adjacency oracle $Adj(v, j)$. The adjacency oracle takes a tree node v and an index $j \in \{0, 1, 2, ..., \Delta - 1\}$. For each value of j in its range, $Adj(v, j)$ either returns *null* or gives a child of v. Each child is given exactly once. Our goal is to visit all of the nodes in T. For concreteness in this section we will describe depth first search (DFS) although the budgeting technique can be applied to breadth first search or other search strategies.

A budgeted DFS is initiated from tree node *start_vertex*, initially set to be the root, and proceeds in a depth first fashion outputting each node as it is generated. We use two stacks, *stack_v* to store vertices and *stack_j* to store indices. We also specify an integer budget param-

[1] `http://cgm.cs.mcgill.ca/~avis/doc/tutorial.html`

eter $b \geq 1$ which causes the tree search to be terminated when b nodes have been generated. At this point the most recently generated vertex v and all of the unexplored siblings along the return path from v to *start_vertex* are returned with the flag *unexplored=true*. All previously generated nodes were output with the flag *unexplored=false*. The pseudo-code is shown as Algorithm 1. Note that that *start_vertex* is not output so this should be done in the calling program. We exploit this property later.

Algorithm 1 Budgeted depth first search

```
1: procedure BDFS(start_vertex, Δ, Adj, b )
2:     j ← 0   v ← start_vertex   count ← 0   depth ← 0
3:     repeat
4:         unexplored ← false
5:         while j < Δ and unexplored = false  do
6:             j ← j + 1
7:             push(stack_v, v)
8:             push(stack_j, j)
9:             v ← Adj(v, j)
10:            depth ← depth + 1    count ← count + 1
11:            if count ≥ b then                          ▷ budget is exhausted
12:                unexplored ← true
13:            end if
14:            output (v, unexplored)
15:            if count < b then                          ▷ continue down tree
16:                j ← 0
17:            end if
18:        end while
19:        if depth > 0 then                              ▷ backtrack step
20:            v ← pop(stack_v)
21:            j ← pop(stack_j)
22:            depth ← depth − 1
23:        end if
24:    until depth = 0 and j = Δ
25: end procedure
```

Consider the tree in Figure 1 which has 25 nodes, $\Delta = 5$ and is rooted at vertex 0. For convenience the nodes are numbered 0,1,...,24 in depth first search order but this is in no way essential. If we set the parameter b to be 25 or greater, all nodes are output in this order with *unexplored=false*. Suppose we set $b = 13$. Firstly nodes 1,...,12 are output in order with *unexplored=false*. Then nodes 13,15,16,18,22 are output with *unexplored=true*. On the other hand, if we set $b = 8$ then nodes 1,2,...,7 are output with *unexplored=false* and nodes 8,9,10,11,15,16,18,22 are output with *unexplored=true*.

We will store the unexplored nodes in a list L which, for DFS, would be implemented as a stack. To complete the DFS of the tree we search the subtrees rooted by nodes in L using Algorithm 1 repeatedly. Note that each vertex in L now becomes a *start_vertex* and if a budget constraint b is applied additional nodes may be added to L. If L is empty when Algorithm 1 terminates the tree has been completely output. Since the root nodes are not output we will obtain a duplicate free list of all the nodes of the tree by concatenating outputs. Note that if the order of output is not important, the subtrees rooted by nodes in L can be explored in any order and in parallel. Furthermore, if $b = 1$ then all nodes in T will be returned to L. Alternatively, if $b > |T|$ then L remains empty as the tree is explored with a single call to Algorithm 1.

A successful parallelization of Algorithm 1 depends on controlling the list L. One one hand,

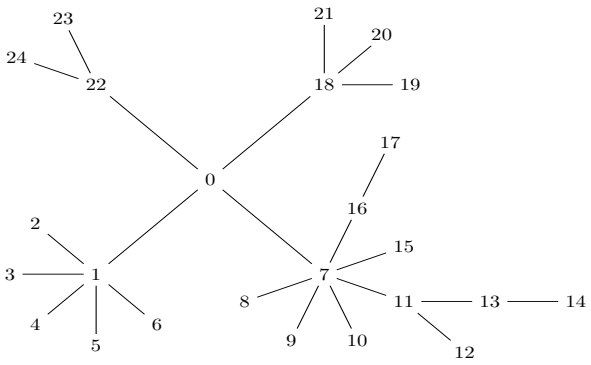

Figure 1: Tree rooted at vertex 0 with 25 nodes and $\Delta = 6$

if the budget parameter b is too small, then L tends to become large causing considerable overhead in restart costs. If b is too large, then L becomes smaller, reducing overhead, but may become empty before the entire tree has been searched. This causes processors to become idle and reduces the speedup that can be obtained. We are interested in how L evolves over time.

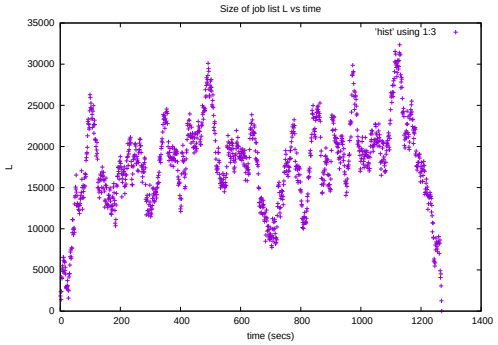

(a) $b = 50,\ R_n = 52,454,166$

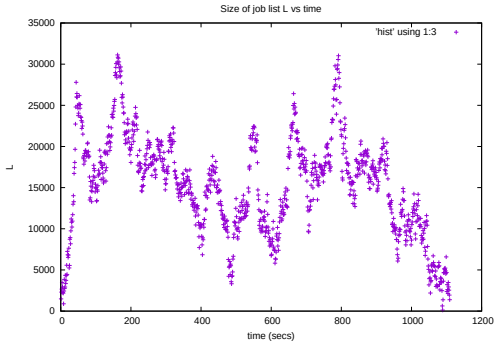

(b) $b = 500,\ R_n = 17,384,235$

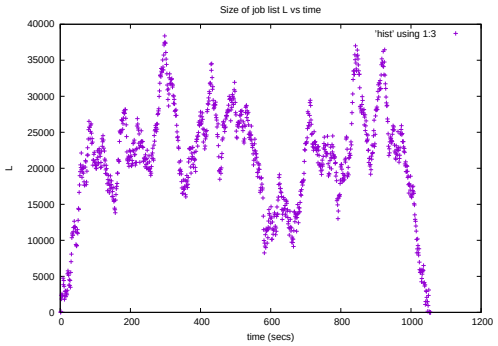

(c) $b = 5000,\ R_n = 5,584,327$

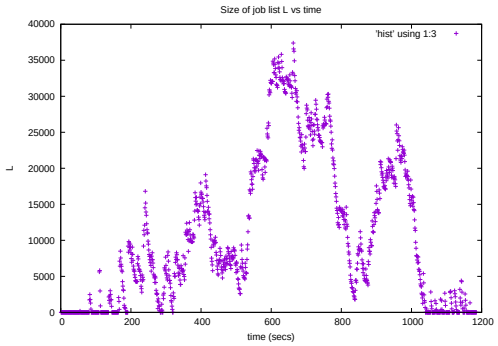

(d) $b = 500,000,\ R_n = 572,507$

Figure 2: Evolution of the job list for T_{10}: budget=b, total restarts=R_n, *mai20*, 16 cores

Let L be the list generated by repeatedly applying Algorithm 1 to a tree T for a given budget b and let R_n denote the total number of nodes added to the list during a complete search of T. We can observe the size of L each time Algorithm 1 is called. For example, again referring to Figure 1, Algorithm 1 is called a total of 6 times if we keep $b = 13$ throughout. $|L_{13}| = 5$ and the size of L_{13} is respectively 0,4,3,2,1,0 for these 6 times.

For a larger example consider Figure 2. This is based on visiting all the nodes of the tree T_{10} which has $\Delta = 10$ and 264,933,315 nodes. The plots were obtained by running Algorithm 1 in parallel on 16 cores using *mts*. The total running time is smallest for $b = 500$ and $b = 5000$. For $b = 50$ the large size of the list L increases the running time. For $b = 500,000$ the list is too small and and becomes empty at the start and end of the run causing cores to be idle. We notice that the size of the list L decreases as the budget is scaled up. In fact by comparing the results for $b = 50, 5000, 500000$ we see that as the budget scales up by a factor of 100, the size of L scales down by roughly a factor of 10, the square root of the scale factor. In the following sections we show that this is typical behaviour for Galton-Watson trees, of which T_{10} is a sample.

Before proceeding we note two important points. Firstly the total size of the job list is independent of the number of workers assigned. Indeed a single worker repeatedly executing Algorithm 1 produces the same number of jobs as a thousand workers in parallel. Secondly, it the job list stays well populated for the duration of the run then workers are never idle. If the jobs on the list are independent, as they are in a Galton-Watson tree, then each worker behaves in the same way in probability. It is therefore sufficient to study the behaviour of a single worker applying Algorithm 1 repeatedly until the job list is empty.

3 The probabilistic model and our main result

A Galton-Watson (or Galton-Watson-Bienaymé) tree (see Athreya and Ney, 1972) is a rooted random ordered tree. Each node independently generates a random number of children drawn from a fixed offspring distribution ξ. The distribution of ξ defines the distribution of T, a random Galton-Watson tree. In what follows, we are mainly interested in critical Galton-Watson trees, i.e., those having $\mathrm{E}\{\xi = 1\}$, and $\mathrm{P}\{\xi = 1\} < 1$). In addition, we assume that the variance of ξ is finite (and hence, nonzero).

Moon (1970) and Meir and Moon (1978) defined the *simply generated trees* as ordered labelled trees of size n that are all equally likely given a certain pattern of labeling for each node of a given degree. The most important examples include the Catalan trees (equiprobable binary trees), the equiprobable k-ary trees, equiprobable unary-binary trees (ordered trees with up to two children), random Motzkin trees, random planted plane trees (equiprobable ordered trees of unlimited degrees) and Cayley trees (equiprobable unordered rooted trees). It turns out that all these trees can be represented as critical Galton-Watson trees conditional on their size, n, a fact first pointed out by Kennedy (1975), and further developed by Kolchin (1980, 1986) and others. So, let T_n be a Galton-Watson tree conditional on its size being n. For example, when ξ is 0 or 2 with probability 1/4, and 1 with probability 1/2, we obtain the uniform binary (Catalan) tree. Uniformly random full binary trees are obtained by setting $\mathrm{P}\{\xi = 0\} = \mathrm{P}\{\xi = 2\} = 1/2$. A uniformly random k-ary tree has its offspring distributed as a binomial $(k, 1/k)$ random variable. A uniform planted plane tree is obtained for the geometric law $\mathrm{P}\{\xi = i\} = 1/2^{i+1}$, $i \geq 0$. When ξ is Poisson of parameter 1, one obtains (the shape of) a random rooted labeled (or Cayley) tree. For ξ uniform on $\{0, 1, 2, \ldots, k\}$, T_n is like a uniform ordered tree with maximal degree of k. All such trees can be dealt with at once in the Galton-Watson framework.

Recall that R_n is the number of restarts generated when Algorithm 1 is used repeatedly to explore a tree T with n nodes using a budget b. Our main result is:

THEOREM 1. *If T_n is a Galton-Watson tree of size n determined by ξ, where $\mathrm{E}\{\xi\} = 1$ and $0 < \mathrm{V}\{\xi\} \stackrel{\text{def}}{=} \sigma^2 < \infty$, then*

$$\frac{R_n}{n} \to \frac{1}{\mu_b}$$

in probability as $n \to \infty$, where

$$\mu_b \stackrel{\text{def}}{=} \mathrm{E}\{\min(|T|, b)\},$$

and T is an unconditional Galton-Watson tree for the same ξ. In addition, as $b \to \infty$,

$$\mu_b \sim \sqrt{\frac{8b}{\pi\sigma^2}}.$$

4 Preliminary results: the random walk view

Let ξ be a random variable representing the number of children of the root in a critical Galton-Watson tree: $\mathrm{E}\{\xi\} = 1$, $\mathrm{P}\{\xi = 0\} > 0$. Set $X = \xi - 1$, and let $X_1, X_2, \ldots$ be a sequence of i.i.d. random variables distributed as X. Define the partial sums $Q(0) = 1$, $Q(t) = 1 + \sum_{i=1}^{t} X_i$. There is a well-known depth first or preorder construction of a random Galton-Watson tree which lends itself well to the study of all properties of randomly selected nodes (see, e.g., Le Gall, 1989, or Aldous, 1991). Nodes in an ordered tree can be encoded with a vector of child numbers. The root corresponds to the empty vector. Its children have encodings $1, 2, \ldots, \xi$. The children of the $j-th$ child of the root are encoded by $(j, 1), (j, 2), \ldots$. A preorder listing of the nodes is nothing but a lexicographic listing of the node vectors.

We can traverse a random Galton-Watson tree by visiting nodes in preorder, starting at the root. To do so, a list V of nodes to be visited is kept, which is initially of size one (V just contains the root). When node u is visited, we consider ξ_u, the number of children of u, and remove u from V. Thus, V increases by $\xi_u - 1$. The next node in lexicographic order is taken from V, and the process continues until V is empty.

We denote by N the size of a random Galton-Watson tree. The size of V after t nodes have been processed is precisely the Q process introduced above, $Q(t)$. Thus, $Q(0) = 1$, and

$$Q(t) = 1 + X_1 + \cdots + X_t,$$

where $X_t = \xi_t - 1$, and indexing of the nodes is by their lexicographic rank. The root has rank one, for example. We have the identity

$$[N = n] = [Q(1) > 0, Q(2) > 0, \ldots, Q(n-1) > 0, Q(n) = 0].$$

The standard circular symmetry argument for random walks (see, e.g., Dwass, 1968) shows that

$$\begin{aligned}
\mathrm{P}\{N = n\} &= \mathrm{P}\{Q(0) = 1, Q(1) > 0, Q(2) > 0, \ldots, Q(n-1) > 0, Q(n) = 0\} \\
&= \frac{1}{n}\mathrm{P}\{Q(n) = 0\} \\
&= \frac{1}{n}\mathrm{P}\{X_1 + \cdots + X_n = -1\}.
\end{aligned}$$

The asymptotics for this probability distribution are well-known, and will be given below. Let $0 < \sigma^2 \stackrel{\text{def}}{=} \mathrm{V}\{\xi\} > 0$ (which implies $\mathrm{P}\{\xi = 1\} < 1$), and let the span d be the greatest common divisor of all $i > 0$ for which $\mathrm{P}\{\xi = i\} > 0$. A Galton-Watson tree with span d can only have sizes that are 1 mod d. From Petrov (1975, p. 197) or Kolchin (1986, p. 16, p. 105), we recall that

$$\lim_{n\to\infty} \sup_{k=1 \bmod d} \left| \sigma\sqrt{n}\,\mathrm{P}\{X_1 + \cdots + X_n = k\} - \frac{d}{\sqrt{2\pi}} e^{-\frac{k^2}{2\sigma^2 n}} \right| = 0,$$

while, clearly, $\mathrm{P}\{X_1 + \cdots + X_n = k\} = 0$ if $k \neq 1 \bmod d$. Thus, along $n \in \mathcal{N}$,

$$\mathrm{P}\{N = n\} \sim \frac{d}{\sigma\sqrt{2\pi}n^{3/2}},$$

and, as $n \to \infty$,

$$\mathrm{P}\{N \geq n\} \sim \sqrt{\frac{2}{\pi n \sigma^2}}.$$

LEMMA 1. *Under the conditions of Theorem 1, as $b \to \infty$,*

$$\mathrm{E}\{\min(N, b)\} \sim \sqrt{\frac{8b}{\pi\sigma^2}}.$$

PROOF. We have

$$\begin{aligned}
\mathrm{E}\{\min(N,b)\} &= \sum_{t=1}^{\infty} t\mathrm{P}\{\min(N,b) = t\} \\
&= \sum_{t=1}^{b} t\mathrm{P}\{N = t\} + b\mathrm{P}\{N > b\} \\
&= \sum_{t=1}^{b} \mathrm{P}\{X_1 + \cdots + X_t = -1\} + b\mathrm{P}\{N > b\} \\
&= \sum_{t \leq b : t \in \mathcal{N}} \frac{d + o(1)}{\sigma\sqrt{2\pi t}} + (1 + o(1))\sqrt{\frac{2b}{\pi\sigma^2}} \\
&\sim \int_1^b \frac{1}{\sigma\sqrt{2\pi t}}\, dt + \sqrt{\frac{2b}{\pi\sigma^2}} \\
&\sim \sqrt{\frac{8b}{\pi\sigma^2}}. \square
\end{aligned}$$

Consider the Q-process conditional on visiting (x, y), with $0 < x < n$ and $y > 0$. Define

$$S(x, y) = \min\{t > x : Q(t) = Q(x) - 1\} - x,$$

where we recall that $Q(t) = 1 + X_1 + \cdots + X_t$. In other words, $S(x, y)$ is the size of the subtree of node x in depth first search order. Note that $S(x, y) = 1 \bmod d$. In addition, since $Q(0) = 1$, we must have $t + Q(t) = 1 \bmod d$, so that the only values (x, y) of interest to us have $x + y = 1 \bmod d$. We consider two positive constants a, c, with $a < 1/2, c < 1$, and let $\mathcal{S}(a, c)$ denote the collection of all random variables $S(x, y)$ with $an \leq x \leq (1 - a)n$, $c\sqrt{n} \leq y \leq (1/c)\sqrt{n}$. Lemma 2 below shows that within $\mathcal{S}(a, c)$, all random variables are close to N, the size of an unconditional Galton-Watson tree. This is an explicit version of a well-known theorem due to Aldous (1991) (see also Janson (2012)), which states that if T_n^* is the subtree of a uniform random node of T_n, then T_n^* tends in distribution to an unconditional Galton-Watson tree T as $n \to \infty$.

LEMMA 2. *For any fixed $0 < a < 1/2$, $0 < c < 1$ and integer $b > 0$,*

$$\lim_{n\to\infty} \sup_{1\le i\le b; i=1 \bmod d} \sup_{(x,y):\begin{cases} S(x,y)\in\mathcal{S}(a,c) \\ x+y=1 \bmod d \end{cases}} \left| \frac{\mathrm{P}\{S(x,y)=i\}}{\mathrm{P}\{N=i\}} - 1 \right| = 0.$$

Furthermore,

$$\lim_{n\to\infty} \sup_{(x,y):\begin{cases} S(x,y)\in\mathcal{S}(a,c) \\ x+y=1 \bmod d \end{cases}} |\mathrm{E}\{\min(b,S(x,y))\} - \mathrm{E}\{\min(b,N)\}| = 0.$$

PROOF. The last part follows from the first one, so we only consider $\mathrm{P}\{S(x,y)=i\}$ for $1 \le i \le b$, $i = 1 \bmod d$. We have, for $x + y = 1 \bmod d$, $n = 1 \bmod d$,

$$\begin{aligned} &\mathrm{P}\{Q(x)=y, Q(t)>0 \text{ for all } x<t<n, Q(n)=0\} \\ &\quad = \mathrm{P}\{Q(x)=y\} \times \frac{1}{n-x}\mathrm{P}\{X_1+\cdots+X_{n-x}=-y\} \\ &\quad = \mathrm{P}\{Q(x)=y\} \times \frac{d+o(1)}{(n-x)\sqrt{2\pi(n-x)\sigma^2}} \exp\left(-\frac{y^2}{2(n-x)\sigma^2}\right). \end{aligned}$$

Also,

$$\begin{aligned} &\mathrm{P}\{Q(x)=y, S(x,y)=i, Q(t)>0 \text{ for all } x+i<t<n, Q(n)=0\} \\ &\quad = \mathrm{P}\{Q(x)=y\} \times \mathrm{P}\{N=i\} \times \frac{1}{n-x-i}\mathrm{P}\{X_1+\cdots+X_{n-x-i}=-(y-1)\} \\ &\quad = \mathrm{P}\{Q(x)=y\} \times \mathrm{P}\{N=i\} \times \frac{d+o(1)}{(n-x-i)\sqrt{2\pi(n-x-i)\sigma^2}} \exp\left(-\frac{(y-1)^2}{2(n-x-i)\sigma^2}\right). \end{aligned}$$

Thus, as

$$\mathrm{P}\{S(x,y)=i\} = \frac{\mathrm{P}\{Q(x)=y, S(x,y)=i, Q(t)>0 \text{ for all } x+i<t<n, Q(n)=0\}}{\mathrm{P}\{Q(x)=y, Q(t)>0 \text{ for all } x<t<n, Q(n)=0\}},$$

we have

$$\begin{aligned} \frac{\mathrm{P}\{S(x,y)=i\}}{\mathrm{P}\{N=i\}} &= (1+o(1))\left(\frac{n-x}{n-x-i}\right)^{3/2} \exp\left(\frac{y^2}{2(n-x)\sigma^2} - \frac{(y-1)^2}{2(n-x-i)\sigma^2}\right) \\ &= (1+o(1)) \exp\left(\frac{y^2}{2(n-x)\sigma^2} - \frac{(y-1)^2}{2(n-x-i)\sigma^2}\right) \\ &= (1+o(1)) \exp\left(\frac{y^2-(y-1)^2}{2(n-x)\sigma^2}\right) \\ &= (1+o(1)) \exp\left(\frac{2y-1}{2(n-x)\sigma^2}\right) \\ &= 1+o(1), \end{aligned}$$

where all the $o(1)$ terms are uniform over all i, x, y within the given ranges. □

The Q process can be viewed as a continuous curve on $[0, n]$ by using linear interpolation between $(t, Q(t))$ and $(t+1, Q(t+1))$. When properly rescaled, it converges in distribution to Brownian excursion. Brownian excursion is Brownian motion $B(t)$ on $[0, 1]$ conditioned to be positive and to take the value 0 at time 1. Alternatively, it is a Brownian bridge process conditioned to be positive. Another representation of a Brownian excursion $e(t)$ due to Paul Lévy (and noted by Ito and McKean, 1974) is in terms of the last time τ_- that Brownian motion

$B(t)$ hits zero before time 1 and the first time τ_+ that Brownian motion $B(t)$ hits zero after time 1:

$$\{e(t) : 0 \le t \le 1\} \stackrel{\mathcal{L}}{=} \left\{ \frac{|B((1-t)\tau_- + t\tau_+)|}{\sqrt{\tau_+ - \tau_-}} : 0 \le t \le 1 \right\}.$$

Since B, and thus e, are continuous processes, we note the following:

(i) $\max_{0\le t\le 1} e(t)$ is a random variable. In fact, it has the theta distribution, as proved by Rényi and Szekeres (1967), Kennedy (1976), Flajolet and Odlyzko (1982) and others.

(ii) For every $\epsilon \in (0, 1/2)$, $\min_{\epsilon\le t\le 1-\epsilon} e(t)$ is a random variable without an atom at zero (for otherwise, it would contradict Lévy's representation).

The convergence of partial sums of i.i.d. zero mean random variables with a finite variance to Brownian motion is well-known. As a consequence, partial sums conditioned on being positive and attaining 0 at the start and end, when suitably rescaled, converge to Brownian excursion. In particular,

$$\left(\frac{Q(tn)}{\sigma\sqrt{n}}, 0 \le t \le 1 \right) \stackrel{\mathcal{L}}{\to} (e(t), 0 \le t \le 1)$$

in the sup norm metric. In other words, there exists a sequence of Brownian excursions $e_1, e_2, \ldots$, such that

$$\sup_{0\le t\le 1} \left| \frac{Q(tn)}{\sigma\sqrt{n}} - e_n(t) \right| \to 0$$

in probability as $n \to \infty$. The work on this convergence of the Q process goes back to Le Gall (1989, 2005), Aldous (1991, 1993), Bennies and Kersting (2000), Marckert and Mokkadem (2003), Duquesne (2003), and others. A consequence of this and remarks (i) and (ii) above is that for given $\epsilon > 0$ and $\delta > 0$, we can find $\theta > 0$ such that with probability at least $1 - \delta$,

$$\theta\sqrt{n} \le \min_{\epsilon n\le t\le (1-\epsilon)n} Q(t) \le \max_{\epsilon n\le t\le (1-\epsilon)n} Q(t) \le \frac{1}{\theta}\sqrt{n}.$$

Call this event $G(\epsilon, \theta)$.

5 Proof of Theorem 1

We are now ready to prove Theorem 1. Let us denote the sizes of the trees explored by the depth first search steps by $N_1, N_2, N_3, \ldots$, where we note that $1 \le N_i \le b$ for all i. The indices of the nodes at which the depth first search operations are started are denoted by $T_1 < T_2 < \cdots$, with $T_1 = 0$, $T_{i+1} = T_i + N_i$ for all i. We have

$$R_n = R = \min\{r > 0 : T_r = n - 1\}.$$

This forces $Q(T_R) = 1$, and thus $Q(T_R + 1) = 0$. In total, the driving stack has held R_n nodes.

By duality,

$$[R \ge r] \equiv [N_1 + \cdots + N_r < n].$$

Recalling the event $G(\epsilon, \theta)$ from the previous section, and denoting by $I(\epsilon)$ the interval of indices $[\epsilon n, r - \epsilon n]$, we have

$$\mathrm{P}\{R \ge r\} \le \mathrm{P}\left\{ \sum_{i\in I(\epsilon)} N_i < n \right\} = \mathrm{P}\left\{ \sum_{i\in I(\epsilon)} (N_i - \mathrm{E}N_i) < n - \sum_{i\in I(\epsilon)} \mathrm{E}N_i \right\}.$$

Note that for $i \in I(\epsilon)$, we have $T_i \geq i \geq \epsilon n$, and $T_i \leq r - \epsilon n$. Also, for any i and j, by simple conditioning, we have

$$\mathrm{E}\left\{(N_i - \mathrm{E}N_i)(N_j - \mathrm{E}N_j)\right\} = 0.$$

By Chebyshev's inequality, and the fact that $\mathrm{V}\{N_i\} \leq b^2$, we have

$$\mathrm{P}\{R \geq r\} \leq \frac{b^2|I(\epsilon)|}{\left(\sum_{i\in I(\epsilon)} \mathrm{E}N_i - n\right)_+^2},$$

where $(y)_+ = \max(u, 0)$. Let E_i be the event

$$\left[\theta\sqrt{n} \leq Q(T_i) \leq \frac{1}{\theta}\sqrt{n}, \epsilon n \leq T_i \leq (1-\epsilon)n\right].$$

. We have for $i \in I(\epsilon)$,

$$\begin{aligned}
\mathrm{E}N_i &\leq \mathrm{E}\{N_i 1_{[G(\epsilon,\theta)]}\} + b\mathrm{P}\{G^c(\epsilon,\theta)\} \\
&\leq \mu_b(1+o(1))\mathrm{P}\{G(\epsilon,\theta)\} + b\delta \\
&\leq \mu_b(1+o(1)) + b\delta,
\end{aligned}$$

where the $o(1)$ refers to the behavior as $n \to \infty$ for fixed b, ϵ, θ. Similarly,

$$\begin{aligned}
\mathrm{E}N_i &\geq \mathrm{E}\{N_i 1_{[E_i]}\} \\
&\geq \mu_b(1+o(1))\mathrm{P}\{E_i\} \\
&\geq \mu_b(1+o(1))(1 - \mathrm{P}\{G^c(\epsilon,\theta)\}) \\
&\geq \mu_b(1+o(1))(1-\delta).
\end{aligned}$$

Therefore,

$$\begin{aligned}
\mu_b(1+o(1))(1-\delta)(r-2\epsilon n) &\leq \mu_b(1+o(1))(1-\delta)|I(\epsilon)| \\
&\leq \sum_{i\in I(\epsilon)} \mathrm{E}N_i \\
&\leq \left(\mu_b(1+o(1)) + b\delta\right)|I(\epsilon)|. \\
&\leq \left(\mu_b(1+o(1)) + b\delta\right) r.
\end{aligned}$$

For n large enough,

$$\sum_{i\in I(\epsilon)} \mathrm{E}N_i - n \geq \mu_b(1-2\delta)(r-2\epsilon n) - n.$$

To show that for any small $\gamma > 0$, $\mathrm{P}\{R \geq n(1+\gamma)/\mu_b\} \to 0$, we choose first $\epsilon = \gamma/(4\mu_b)$, then $\delta = \gamma/(4+4\gamma)$, and then $\theta > 0$ so small that $\mathrm{P}\{G^c(\epsilon,\theta)\} \leq \delta$. With those choices,

$$\begin{aligned}
\mu_b(1-2\delta)(\mu_b^{-1}(n+n\gamma) - 2\epsilon n) - n &= n(1-2\delta)(1+\gamma/2) - n \\
&= n\frac{\gamma^2}{4+4\gamma} \\
&\stackrel{\mathrm{def}}{=} n\gamma'.
\end{aligned}$$

So, for n large enough, by Chebyshev's inequality invoked above,

$$\mathrm{P}\{R \geq n(1+\gamma)/\mu_b\} \leq \frac{b^2 r}{(n\gamma')^2} \to 0.$$

The other side is dealt with in the same manner. First,

$$\mathrm{P}\{R < r\} \leq \mathrm{P}\left\{\sum_{i\in I(\epsilon)} N_i \geq n\right\}$$

$$= \mathrm{P}\left\{ \sum_{i \in I(\epsilon)} (N_i - \mathrm{E}N_i) \geq n - \sum_{i \in I(\epsilon)} \mathrm{E}N_i \right\}$$
$$\leq \frac{b^2 r}{\left(n - \sum_{i \in I(\epsilon)} \mathrm{E}N_i \right)_+^2}.$$

Note that for n large enough,

$$n - \sum_{i \in I(\epsilon)} \mathrm{E}N_i \geq n - (\mu_b + 2b\delta)\, r.$$

To show that for any small $\gamma > 0$, $\mathrm{P}\{R < n(1-\gamma)/\mu_b\} \to 0$, we choose first $\delta = \gamma\mu_b/(2b)$, and then $\theta > 0$ so small that $\mathrm{P}\{G^c(\epsilon, \theta)\} \leq \delta$. With those choices,

$$n - (\mu_b + 2b\delta)\, \frac{n(1-\gamma)}{\mu_b} = n\gamma^2.$$

So, for n large enough, by Chebyshev's inequality invoked above,

$$\mathrm{P}\{R < n(1-\gamma)/\mu_b\} \leq \frac{b^2 r}{n^2 \gamma^4} \to 0.$$

6 Empirical results and discussion

In this section we start by giving some test results to show the accuracy of Theorem 1. Except for $\Delta = 2$ we considered critical Galton-Watson trees with $p_i = 1/i\Delta, i = 1, 2, ..., \Delta$ and $p_0 = 1 - \sum_{i=1}^{\Delta} p_i$. For this family of trees we have that $\sigma^2 = (\Delta - 1)/2$. For $\Delta = 2$ we used the distribution $p = (1/3, 1/3, 1/3)$ which has $\sigma^2 = 2/3$. For each $\Delta = 2, 3, 5, 10, 20, 40$ we generated random Galton-Watson trees until we found one with at least 10^8 nodes. In Table 1 we show the results of searching these trees with various budgets b. As before we denote the number of restarts by R_n. According to Theorem 1, $R_n/\sigma n$ should asymptotically approach $\sqrt{\pi/8b}$ as b gets large. The results in Table 2 show good empirical agreement with this asymptotic bound even for small values of b.

Budget	R_n					
b	T_2	T_3	T_5	T_{10}	T_{20}	T_{40}
50	36164525	70861643	51881667	52454166	118688002	51764031
500	10656708	21596567	16331233	17384235	40853820	18621338
5000	3293174	6764700	5181648	5584327	13295349	6111917
50000	1029919	2121621	1643866	1758032	4242028	1994306
500000	321993	675660	520932	572507	1340837	620958
n	440480727	722813312	393789282	264933315	434056761	145049024
seed	1486862923	2053907278	2018277783	1691015410	1992061106	2121539929

Table 1: Test results

Budget	$R_n/\sigma n$						Estimate
b	T_2	T_3	T_5	T_{10}	T_{20}	T_{40}	$\sqrt{\pi/8b}$
50	.10055	.09804	.09316	.09333	.08872	.08082	.08862
500	.02963	.02988	.02933	.03093	.03054	.02907	.02802
5000	.00916	.00936	.00936	.00994	.00994	.00954	.00886
50000	.00286	.00294	.00295	.00313	.00317	.00311	.00280
500000	.00090	.00093	.00094	.00102	.00100	.00097	.00089

Table 2: Accuracy of Theorem 1 estimate

We now discuss how results like Theorem 1 are useful in practice. Recall from the introduction that budgeting causes two competing forms of overhead: restart cost and idle time. Let us model a unit of computation time by a single node evaluation in Algorithm 1. We can model the restart cost for a job as a constant r units of time, i.e., r node evaluations. For example a node evaluation in *mplrs* involves pivoting an m by n matrix. The restart cost in this case is approximately $r = n$, as the input matrix has to be pivoted up to n times to reach a restart cobasis. Note that the restart cost is born by an individual worker and so the total restart cost is independent of the number of workers. As a proportion of the total work required it is precisely the ratio R_n/n that is estimated in Theorem 1.

If the job list becomes empty then any worker returning for work remains idle until some new jobs arrive. The cost is the number of idle workers for each time unit when the list is empty. Unlike restart costs, these costs scale up as the number of workers increases and can greatly reduce the parallelization efficiency achieved. To minimize these costs the job list should always be well populated.

In practice the budget is not fixed for the duration of a run. The master monitors the size of job list and if it drops too low scales the budget down and if it gets too large scales the budget up. Theorem 1 tells us how this scaling will work for Galton-Watson trees: scaling down (up) by a factor of 100 increases (decreases) the work returned on average by a factor of 10. The dependence of the job list size on the square root of the budget explains why relatively small changes in the budget do not greatly effect performance. This can be observed in Figure 2 for example.

The empirical results obtained by Avis and Jordan (2015, 2016) come from a variety of practical applications. Whilst these trees are not generated according to the Galton-Watson model, the behaviour observed is surprisingly similar to that predicted by Theorem 1. It will be interesting to see if similar theoretical results hold for other classes of trees.

7 Acknowledgements

We gratefully acknowledge helpful conversations with Louigi Addario-Berry and Charles Jordan. The research of Avis was supported by a JSPS Kakenhi Grant and a Grant-in-Aid for Scientific Research on Innovative Areas, 'Exploring the Limits of Computation (ELC)'. The research of Devroye was supported by the Natural Sciences and Engineering Research Council of Canada.

8 References

D. Aldous, "The continuum random tree. II. An overview," in: *Stochastic Analysis (Durham, 1990)*, vol. 167, pp. 23–70, Cambridge Univ. Press, Cambridge, 1991.

D. Aldous, "The continuum random tree. I," *The Annals of Probability*, vol. 19, pp. 1–28, 1991.

D. Aldous, "Asymptotic fringe distributions for general families of random trees," *The Annals of Applied Probability*, vol. 1, pp. 228–266, 1991.

D. Aldous, "The continuum random tree. III," *The Annals of Probability*, vol. 21, pp. 248–289, 1993.

D. Aldous and J. Pitman, "Tree-valued Markov chains derived from Galton-Watson processes," *Annals of the Institute Henri Poincaré*, vol. 34, pp. 637–686, 1998.

K. B. Athreya and P. E. Ney, *Branching Processes*, Springer Verlag, Berlin, 1972.

D. Avis and C. Jordan, "*mplrs*: a scaleable parallel vertex/facet enumeration code," arXiv:1511.06487, 2015.

D. Avis and C. Jordan, "A parallel framework for reverse search using *mts*," arXiv:1610.07735, 2016.

J. Bennies and G. Kersting, "A random walk approach to Galton-Watson trees," *Journal of Theoretical Probability*, vol. 13, pp. 777–803, 2000.

N. Blumofe and C. Leiserson, "Scheduling multithreaded computations by work stealing," *Journal of the ACM*, vol. 46, pp. 720–748, 1999.

T. Duquesne, "A limit theorem for the contour process of conditioned Galton-Watson trees. ," *Annals of Probability*, vol. 31, pp. 996–1027., 2003.

M. Dwass, "The total progeny in a branching process," *Journal of Applied Probability*, vol. 6, pp. 682–686, 1969.

P. Flajolet and A. Odlyzko, "The average height of binary trees and other simple trees," *Journal of Computer and System Sciences*, vol. 25, pp. 171–213, 1982.

P. Flajolet and R. Sedgewick, *Analytic Combinatorics*, Cambridge University Press, Cambridge, 2008.

I. A. Ibragimov, "On the accuracy of the approximation of distribution functions of sums of independent random variables by the normal distribution," *Theory of Probability and its Applications*, vol. 11, pp. 559–580, 1966.

K. Ito and H. P. McKean, *Diffusion Processes and their Sample Paths*, Springer-Verlag, Berlin, 1974.

S. Janson, "Simply generated trees, conditioned Galton-Watson trees, random allocations and condensation," *Probability Surveys*, vol. 9, pp. 103–252, 2012.

S. Janson, "Asymptotic normality of fringe subtrees and additive functionals in conditioned Galton-Watson trees," *Random Structures and Algorithms*, vol. 48, pp. 57–101, 2016.

D. P. Kennedy, "The Galton-Watson process conditioned on the total progeny," *Journal of Applied Probability*, vol. 12, pp. 800–806, 1975.

D. P. Kennedy, "The distribution of the maximum Brownian excursion," *Journal of Applied Probability*, vol. 13, pp. 371–376, 1976.

V. F. Kolchin, "Branching Processes and random trees," in: *Problems in Cybernetics, Combinatorial Analysis and Graph Theory (in Russian)*, pp. 85–97, Nauka, Moscow, 1980.

V. F. Kolchin, *Random Mappings*, Optimization Software Inc, New York, 1986.

J.-F. Le Gall, "Marches aléatoires, mouvement Brownien et processus de branchement," in: *Séminaire de Probabilités XXIII*, edited by J. Azéma, P. A. Meyer and M. Yor, vol. 1372, pp. 258–274, Lecture Notes in Mathematics, Springer-Verlag, Berlin, 1989.

J.-F. Le Gall, "Random trees and applications," *Probability Surveys*, vol. 2, pp. 245–311, 2005.

J. F. Marckert and A. Mokkadem, "The depth first processes of Galton-Watson trees converge to the same Brownian excursion," *Annals of Probability*, vol. 31, pp. 1655–1678, 2003.

T. Mattson, B. Saunders and B. Massingill, *Patterns for Parallel Programming*, Addison-Wesley, 2004.

A. Meir and J. W. Moon, "On the altitude of nodes in random trees," *Canadian Journal of Mathematics*, vol. 30, pp. 997–1015, 1978.

J. W. Moon, *Counting Labelled Trees*, Canadian Mathematical Congress, Montreal, 1970.

J. Neveu and J. W. Pitman, "The branching process in a Brownian excursion," in: *Séminaire de Probabilités XXIII*, edited by J. Azéma, P. A. Meyer and M. Yor, vol. 1372, pp. 248–257, Lecture Notes in Mathematics, Springer-Verlag, Berlin, 1989.

V. V. Petrov, *Sums of Independent Random Variables*, Springer-Verlag, Berlin, 1975.

A. Rényi and G. Szekeres, "On the height of trees," *Journal of the Australian Mathematical Society*, vol. 7, pp. 497–507, 1967.

C. Xu and F, Lau, *Load Balancing in Parallel Computers: Theory and Practice*, Kluwer Academic Publishers, Norwell, MA, 1997.